Adviser FAQ — agentdatasync

1. Does this replace my CRM? No. Your CRM (Wealthbox, Redtail, Salesforce, or whatever you run) stays. agentdatasync cleans a copy of its data so your AI assistants can read it safely, and your firm approves every cleanup change before anything is applied.

2. Which CRMs does it work with? It's built to intake standard exports from Wealthbox, Redtail, and Salesforce first, with other systems on a case-by-case basis at onboarding. The intake mapping is done jointly with your firm, so oddball fields get handled deliberately rather than guessed at.

3. Which AI tools can connect to it? Any agent tool that speaks the Model Context Protocol — the open standard for connecting AI tools to data. That includes tools like Claude and ChatGPT integrations, Microsoft Copilot, and in-house tools your firm builds. It's agent-agnostic by design: you choose the tools, we supply the governed data layer.

4. What does the CCO get out of this? Three things: (1) a **Security & Audit whitepaper** mapping every control to Reg S-P and Rule 204-2, with an honest column showing what stays the firm's responsibility; (2) an **audit log** of every agent query, exportable for examinations; and (3) a **monthly review checklist** so oversight is a repeatable 20-minute task, not a research project. The CCO also approves the permissions matrix before any agent connects.

5. What happens if there's a breach? Two clocks. agentdatasync must **notify your firm within 72 hours** of becoming aware of a breach on our side (that commitment is in the service agreement — Reg S-P effectively requires firms to demand it of their vendors). Your firm then assesses and, where required, **notifies affected clients within 30 days**. Per-agent credentials mean a compromised assistant can be cut off individually, immediately.

6. Can agents change anything in my data? No. The tools we expose are read-only — search and retrieval only. There is no write, delete, or action tool for an agent to misuse. Humans make changes; agents read.

7. Can it move money or place trades? Never. It cannot move money, execute trades, generate advice, or communicate with clients. That boundary is architectural — it's not a setting that could be toggled on.

8. Is client data used to train AI models? No. Queries pass through the system; nothing is retained for training any model — ours, the agent vendor's, or anyone's.

9. Can I see what an agent asked about my client? Yes — that's what the audit log is for. Every query records the timestamp, the agent identity, the question as asked, and the data classes returned. The CCO reviews it monthly and exports it to the firm's records.

10. What happens when an employee leaves or we offboard the product? Per-agent credentials make this surgical: the departing employee's agents get their credentials revoked without disrupting anyone else. On full offboarding, agentdatasync securely deletes the cleaned dataset and revokes all credentials, then confirms deletion in writing. The audit log is retained per the firm's recordkeeping instructions.

11. How does it handle sensitive fields like Social Security numbers and account numbers? Field-level protection is built in: account numbers are returned **masked (last 4 only)** by default, dates of birth are **year-only** unless the CCO justifies full access in the permissions matrix, and **tax IDs are never returned in full to any agent**. These aren't optional features — they're the shipping defaults.

12. What does it cost? *Proposal — founder confirms. Do not quote these as final.* The planned model is a flat **per-firm platform fee** (not per-seat, so it never reads as a CRM price comparison), plus the **Audit tier** — the examiner-ready audit log, monthly review tooling, and retention support — as the paid upsell. Exact numbers are set by the founder; ask us and we'll confirm in writing.